

ARM Processor CortexTM-A15 MPCore - NEON

Product Revision r0p4

Software Developers Errata Notice

Non-Confidential - Draft



Software Developers Errata Notice

Copyright © 2012 ARM. All rights reserved.

Non-Confidential Proprietary Notice

This document is protected by copyright and the practice or implementation of the information herein may be protected by one or more patents or pending applications. No part of this document may be reproduced in any form by any means without the express prior written permission of ARM. **No license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document.**

This document is **Non-Confidential** but any disclosure by you is subject to you providing the recipient the conditions set out in this notice and procuring the acceptance by the recipient of the conditions set out in this notice.

Your access to the information in this document is conditional upon your acceptance that you will not use, permit or procure others to use the information for the purposes of determining whether implementations infringe your rights or the rights of any third parties.

Unless otherwise stated in the terms of the Agreement, this document is provided "as is". ARM makes no representations or warranties, either express or implied, included but not limited to, warranties of merchantability, fitness for a particular purpose, or non-infringement, that the content of this document is suitable for any particular purpose or that any practice or implementation of the contents of the document will not infringe any third party patents, copyrights, trade secrets, or other rights. Further, ARM makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of such third party patents, copyrights, trade secrets, or other rights.

This document may include technical inaccuracies or typographical errors.

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT LOSS, LOST REVENUE, LOST PROFITS OR DATA, SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL OR PUNITIVE DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF OR RELATED TO ANY FURNISHING, PRACTICING, MODIFYING OR ANY USE OF THIS DOCUMENT, EVEN IF ARM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Words and logos marked with ® or TM are registered trademarks or trademarks, respectively, of ARM Limited. Other brands and names mentioned herein may be the trademarks of their respective owners. Unless otherwise stated in the terms of the Agreement, you will not use or permit others to use any trademark of ARM Limited.

This document consists solely of commercial items. You shall be responsible for ensuring that any use, duplication or disclosure of this document complies fully with any relevant export laws and regulations to assure that this document or any portion thereof is not exported, directly or indirectly, in violation of such export laws.

In this document, where the term ARM is used to refer to the company it means "ARM or any of its subsidiaries as appropriate".

Copyright © 2012 ARM Limited
110 Fulbourn Road, Cambridge, England CB1 9NJ. All rights reserved.

Web Address

<http://www.arm.com>

Feedback on content

If you have any comments on content, then send an e-mail to errata@arm.com. Give:

- the document title
- the document number, ARM-EPM-025235
- the page numbers to which your comments apply
- a concise explanation of your comments.

ARM also welcomes general suggestions for additions and improvements.

Release Information

Errata are listed in this section if they are new to the document, or marked as “updated” if there has been any change to the erratum text in Chapter 2. Fixed errata are not shown as updated unless the erratum text has changed. The summary table in section 2.2 identifies errata affect the r0p4 product revision.

04-Sept 2012: Changes in Document v0.3

First issue of the r0p4 Software Developers Errata Notice

DRAFT

Contents

CHAPTER 1.	6
INTRODUCTION	6
1.1. Scope of this document	6
1.2. Categorization of errata	6
CHAPTER 2.	7
ERRATA DESCRIPTIONS	7
2.1. Product Revision Status	7
2.2. Revisions Affected	7
2.3. Category A	9
2.4. Category A (Rare)	9
773769: Large data RAM latencies can lead to rare data corruption	9
775619: L2 may deadlock or corrupt data if all ways in a single set become blocked from replacement.....	10
2.5. Category B	11
766170: HCR.FB doesn't upgrade ICIALLU to ICIALUIS.....	11
766171: TLBINV followed by changing VTTBR VMID value could cause incorrect TLB invalidation	12
766173: Prefetch hit by WB-No Allocate load can incorrectly set inclusion bit	13
766421: Strongly-Ordered/Device load or NC LDREX could return incorrect data	14
766422: Thumb store translation fault to Hypervisor may not have correct HSR Rt value	16
769270: L2 Tag Latency = 3 should not be used in a system configured with Tag Slice	17
773022: Incorrect instructions may be executed from loop buffer	18
774569: Watchpoint may not be taken on second half of unaligned ld/st crossing a 64-byte aligned boundary ..	19
774769: Data corruption may occur with store streaming in a system	20
777770: ESR incorrect for AdvSID HCPTR trapped inst in Hyp mode	21
784420: Speculative instruction fetches with MMU disabled might not comply with architectural requirements	22
784477: CTIINTACK register needs clearing each time it is set.....	23
785769: Undefined exception is not generated for LDC/STC instructions which access DCC registers in User mode.....	24
2.6. Category B (Rare)	24
773319: Unaligned page boundary crossing load hit by TLB invalidate can stall until next interrupt	24
2.7. Category C	25
763126: Three processor exclusive access livelock	25
766174: ID_PFR0 Jazelle extension support incorrectly encoded.....	27
766419: Benign op instead of UNDEFINED in certain unused opcodes.....	28
768529: Debug version of ID_PFR0 Jazelle extension support incorrectly encoded	29
768532: Mapping the GIC memory mapped region as cacheable can result in unpredictable behavior or a deadlock.....	30
770320: Single bit ECC error can cause cache maintenance operation to violate memory ordering	31
773023: Order may not be maintained between Strongly Ordered memory requests and Device memory requests on ACE/AXI	32
774570: Fault Status bit in register DBGDSCR is implemented as sticky.....	33
774571: Sampling registers DBGCIDSR and DBGVIDSR can return incorrect value	34
775622: Debug version of Cache Type Register (CTR) IminLine field are incorrectly encoded	35
777769: ICache parity error may not be corrected for NC code	36

777771:	Ordering of read accesses to the same memory location may not be ensured in the case of an unaligned load	37
777772:	Device LDM may stall if memory type changed asynchronously from Device to Normal	38
777774:	DCCMVAC/DCCSW colliding with ReadOnce from ACE could cause data corruption	39
780121:	PTM might not acknowledge a trace flush request when cpu is in WFI.	41
784419:	An unaligned load crossing a 4K boundary between SO/Dev memory and WB memory can stall the core until the next interrupt.....	42
784469:	CTI Authentication Status register is incorrect.....	43

DRAFT

Chapter 1.

Introduction

This chapter introduces the errata notice for the ARM Cortex-A15 processor.

1.1. Scope of this document

This document describes errata categorized by level of severity. Each description includes:

- the current status of the defect
- where the implementation deviates from the specification and the conditions under which erroneous behavior occurs
- the implications of the erratum with respect to typical applications
- the application and limitations of a ‘work-around’ where possible

This document describes errata that may impact anyone who is developing software that will run on implementations of this ARM product.

1.2. Categorization of errata

Errata recorded in this document are split into the following levels of severity:

Table 1 **Categorization of errata**

Errata Type	Definition
Category A	A critical error. No workaround is available or workarounds are impactful. The error is likely to be common for many systems and applications.
Category A(rare)	A critical error. No workaround is available or workarounds are impactful. The error is likely to be rare for most systems and applications. Rare is determined by analysis, verification and usage.
Category B	A significant error or a critical error with an acceptable workaround. The error is likely to be common for many systems and applications.
Category B(rare)	A significant error or a critical error with an acceptable workaround. The error is likely to be rare for most systems and applications. Rare is determined by analysis, verification and usage.
Category C	A minor error.

Chapter 2.

Errata Descriptions

2.1. Product Revision Status

The *mpn* identifier indicates the revision status of the product described in this book, where:

- rn*** Identifies the major revision of the product.
- pn*** Identifies the minor revision or modification status of the product.

2.2. Revisions Affected

Table 2 below lists the product revisions affected by each erratum. A cell marked with **X** indicates that the erratum affects the revision shown at the top of that column.

This document includes errata that affect revision r0p4 only.

Refer to the reference material supplied with your product to identify the revision of the IP.

Table 2 Revisions Affected

ID	Cat	Rare	Summary of Erratum	r0p4
773769	CatA	Rare	Large data RAM latencies can lead to rare data corruption	X
775619	CatA	Rare	L2 may deadlock or corrupt data if all ways in a single set become blocked from replacement	X
766170	CatB		HCR.FB doesn't upgrade ICIALLU to ICIALLUIS	X
766171	CatB		TLBINV followed by changing VTTBR VMID value could cause incorrect TLB invalidation	X
766173	CatB		Prefetch hit by WB-No Allocate load can incorrectly set inclusion bit	X
766421	CatB		Strongly-Ordered/Device load or NC LDREX could return incorrect data	X
766422	CatB		Thumb store translation fault to Hypervisor may not have correct HSR Rt value	X
769270	CatB		L2 Tag Latency = 3 should not be used in a system configured with Tag Slice	X
773022	CatB		Incorrect instructions may be executed from loop buffer	X
774569	CatB		Watchpoint may not be taken on second half of unaligned ld/st crossing a 64-byte aligned boundary	X
774769	CatB		Data corruption may occur with store streaming in a system	X
777770	CatB		ESR incorrect for AdvSID HCPTR trapped inst in Hyp mode	X
784420	CatB		Speculative instruction fetches with MMU disabled might not comply with architectural requirements	X
784477	CatB		CTIINTACK register needs clearing each time it is set	X
785769	CatB		Undefined exception is not generated for LDC/STC instructions which access DCC registers in User mode	X
773319	CatB	Rare	Unaligned page boundary crossing load hit by TLB invalidate can stall until next interrupt	X
763126	CatC		Three processor exclusive access livelock	X

ID	Cat	Rare	Summary of Erratum	Impact
766174	CatC		ID_PFR0 Jazelle extension support incorrectly encoded	X
766419	CatC		Benign op instead of UNDEFINED in certain unused opcodes	X
768529	CatC		Debug version of ID_PFR0 Jazelle extension support incorrectly encoded	X
768532	CatC		Mapping the GIC memory mapped region as cacheable can result in unpredictable behavior or a deadlock	X
770320	CatC		Single bit ECC error can cause cache maintenance operation to violate memory ordering	X
773023	CatC		Order may not be maintained between Strongly Ordered memory requests and Device memory requests on ACE/AXI	X
774570	CatC		Fault Status bit in register DBGDSCR is implemented as sticky	X
774571	CatC		Sampling registers DBGCIDSR and DBGVIDSR can return incorrect value	X
775622	CatC		Debug version of Cache Type Register (CTR) IminLine field are incorrectly encoded	X
777769	CatC		ICache parity error may not be corrected for NC code	X
777771	CatC		Ordering of read accesses to the same memory location may not be ensured in the case of an unaligned load	X
777772	CatC		Device LDM may stall if memory type changed asynchronously from Device to Normal	X
777774	CatC		DCCMVAC/DCCSW colliding with ReadOnce from ACE could cause data corruption	X
780121	CatC		PTM might not acknowledge a trace flush request when cpu is in WFI.	X
784419	CatC		An unaligned load crossing a 4K boundary between SO/Dev memory and WB memory can stall the core until the next interrupt	X
784469	CatC		CTI Authentication Status register is incorrect	X

2.3. Category A

2.4. Category A (Rare)

773769: Large data RAM latencies can lead to rare data corruption

Category A Rare

Products Affected: Cortex-A15 MP Core -NEON.

Present in: r0p4

Description

In a system with ACP in use, with long L2 data RAM latencies (4-8 cycles), or with an L2 data RAM slice configured, it is possible that a rare collision between non-cacheable stores and L1 data cache evictions can lead to data corruption.

Note: This erratum matches bug #5269 in the ARM internal Jira database.

Configurations Affected

System has one or more devices connected to the A15 ACP port

Conditions

- 1) The L2 Data RAM latency is programmed to 4 or more cycles, OR ACP is in use, OR A15 is configured with one or more Data RAM slices
- 2) Multiple stores are sent from the L1 data memory to the L2 cache
- 3) Memory type of the stores: strongly ordered, device; normal inner-NC, inner-WT, or inner-WBNoAllocate
- 4) Multiple evictions from the L1 occur intermixed with the non-cacheable stores
- 5) The evictions and stores are stalled and hazarded in an extremely rare manner

Implications

Incorrect data will be written to the L2 cache or to memory.

A typical 2-CPU A15 implementation will not use a data slice, and will have a data RAM latency of 3 or less. If ACP is not used, that implementation will not be affected.

If an A15 implementation is using ACP, is configured with a data RAM slice, or is using a slower data RAM, in very rare circumstances data corruption could occur. This issue has been reproduced a single time in a stressful environment with a data RAM latency of 8 cycles. The required boundary conditions become less likely to align with lower latency data RAMs. The issues has not been reproduced in other configurations, but can't be ruled out except in the configurations described above.

The default reset value of the data RAM latency is 8 cycles, and would then be programmed by boot code to the faster value (typically 3-5 cycles) before turning on the MMU or enabling the cache.

Workaround

If a Data RAM slice is configured in the A15, there is no workaround.

If the system uses ACP, ACP masters must not issue memory requests to A15.

Data RAM latency in the L2 should be programmed to 2 or 3 cycles. 4, 5, 6, 7, and 8 cycle data RAMs should never be used. Because the reset value of the L2CTLR[2:0] is 3'b000 (corresponding to 8 cycle data RAM latency), after reset and before the MMU is enabled or the SCTLR.c bit is set, the L2CTLR register data ram latency bits (L2CTLR[2:0]) should be programmed to 3'b001 or 3'b010, corresponding to data RAM latencies of 2 or 3 cycles respectively.

775619: L2 may deadlock or corrupt data if all ways in a single set become blocked from replacement**Category A Rare****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

A15 has bits in each cache line in the Level 2 Cache (L2) that can block the line from being replaced. These bits are set to indicate a line is currently held in an L1 cache and must be retained, or may be temporarily set if there is a multi-pass request in flight to that line.

It is possible for all 16 lines in a given L2 set to be 'locked'. If this occurs, requests that need to replace a line in the L2 should detect a hazard and retry, waiting for a line to be released. However, there are two scenarios in which all 16 ways becoming locked can cause an issue.

If all 16 ways in a given set are locked and an ACP write-allocate request comes, the ACP request will not hazard correctly. It will replace a locked line, putting the L2 cache in an inconsistent state. This could lead to data corruption to the locked line that was replaced.

If all 16 ways in a given set are locked a cache line fill that needs to replace a line in that set will continuously hazard and restart. In a rare set of arbitration conditions, it is possible that that request will block the other requests that will eventually release the locked/inclusive lines. This will lead to a deadlock.

Note: This erratum matches bug #5298 and 5302 in the ARM internal Jira database.

Configurations Affected

- Cat A Rare for 3 or 4 CPU A15 configurations
- Cat B Rare for 1 or 2 CPU A15 configurations

Conditions

- 1) All ways in a given L2 set have their lock or inclusion bit set
- 2) An ACP write allocate request is made to that set OR another cache line fill needs to replace a way in that set

Implications

In an A15 with one or two CPUs, there is an easy zero impact workaround that will prevent all the lock/inclusion bits in a set from ever being set at the same time (see workaround below), completely avoiding the data corruption or deadlock.

In an A15 with three or four CPUs, there is no workaround. A system with no ACP write allocate traffic will avoid the data corruption scenario, but could see a deadlock.

Workaround

This erratum can be completely avoided for an A15 containing only one or two CPUs. The "Force in-order requests to same set/way" bit of the Auxiliary Control Register should be set (ACTLR[23] set to 1'b1) in each CPU. This bit will prevent that CPU from having more than two cache line fills outstanding at a given time to the same set. This will completely prevent the L2 from ever seeing all 16 ways in a set with their inclusion or lock bits set and thus avoid the erratum. Because the current L2 implementation will not process more than two requests at a time to a given L1 data cache set, there is no performance impact to setting this bit.

For an A15 containing three or four CPUs, there is no workaround. Preventing ACP write traffic will prevent the silent data corruption, but there is the possibility of a deadlock.

2.5. Category B

766170: HCR.FB doesn't upgrade ICIALLU to ICIALLUIS

Category B

Products Affected: Cortex-A15 MP Core -NEON.

Present in: r0p4

Description

When the HCR.FB bit is set, instruction cache invalidate all commands (ICIALLU) must be upgraded to instruction cache invalidate all Inner Shared (ICIALLUIS). Instead, A15 executes the ICIALLU command. This means that the instruction cache local to that processor will be invalidated, but none of the instruction caches in the other processors in the A15 cluster will be invalidated.

Note: This erratum matches bugs #5113 in the ARM internal Jira database.

Conditions

- 1) HCR.FB = 1
- 2) Execute ICIALLU

Implications

This feature would be used by a hypervisor to allow a Guest-OS that is not MP aware to be able to migrate between multiple processors in an A15 cluster. If the non-MP aware Guest-OS modifies instruction memory (loading instruction pages or doing self-modifying code in a JIT) and follows it with an ICIALLU to clean the instruction cache, the instruction caches of the other processors would not be invalidated. Were that Guest-OS to migrate to another processor it could find stale copies of instruction memory in the instruction cache. By upgrading the ICIALLU to ICIALLUIS, the hypervisor should have avoided this problem, but the erratum means that won't work

Workaround

When the hypervisor is migrating a non-MP aware OS from processor A to processor B within the A15 cluster, the hypervisor must execute an ICIALLU on processor B, or an ICIALLUIS on any processor. This will prevent the non-MP aware OS from seeing stale data in the icache.

766171: TLBINV followed by changing VTTBR VMID value could cause incorrect TLB invalidation**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

When a TLB invalidate command is executed in non-secure mode, the VMID is often part of the matching criteria used. This VMID should be the current VMID of the executing processor at the time the TLB invalidate command is executed.

However, if a processor executes a TLB Invalidate command followed quickly by a change to the VMID and an ISB to force completion of that VMID change, the TLB Invalidate command may incorrectly use the updated VMID value to qualify its invalidations, rather than the original VMID. TLB entries that should have been invalidated by the command may be left valid in the TLBs.

Note: This erratum matches bugs #5119 in the ARM internal Jira database.

Conditions

- 1) A TLB Invalidate command is executed in non-secure state (Kernel or Hypervisor)
- 2) The Virtualization Translation Table Base Register (VTTBR) is written changing the VMID (VTTBR[55:48])
- 3) An ISB instruction is executed before the TLB invalidate has completed execution

Implications

If this erratum is hit, the TLB will not be correctly invalidated. Incorrect translations could occur.

In general, it would be difficult (but not impossible) for this erratum to affect TLB invalidate operations executed in non-secure Kernel mode. This is because NS-Kernel can't change the VMID and it is very unlikely that TLB operations would not have had time to sample the correct VMID value before a mode change to hypervisor and then a VMID change was completed.

The code most vulnerable to this erratum would be the hypervisor TLB maintenance code. It is possible that the hypervisor would choose to execute one or more TLB entries associated with one VMID, quickly change the VMID, and execute one or more TLB entries associated with another VMID.

Workaround

After any TLB invalidate command and before the next VMID change, execute either a DMB or a DSB instruction. Either of these barriers will force all prior TLB invalidate operations to complete to the point where they will no longer be affected by VMID changes.

The simplest approach would be to find all code that adjusts the VTTBR and place a DMB before the write to the VTTBR.

766173: Prefetch hit by WB-No Allocate load can incorrectly set inclusion bit**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

The hardware prefetcher initiates a read of a write-back read-write-allocate cache line. One of the processors then issues a write-back no-allocate read request to the same line. This hits the outstanding prefetch and merges with it. This request must now recognize that, upon returning data to the requesting processor, the line will not be allocated into the L1 data cache. In the erratum case, however, the L2 inclusion bit (indicating the line is being held in an L1 data cache) is set incorrectly. This prevents it being replaced until the L1 data cache evicts that line. Because the L1 data cache does not have that line, that bit will never be cleared during normal operation and the line can never be replaced.

This will not lead to an immediate problem (other than a negligible loss of performance due to that cache set being 15 way instead of 16 way). However, over a very long period of time it would be possible for this rare boundary condition to occur multiple times for the same set. If it happened enough times for the same set (minimum eight), the A15 could deadlock.

Write-back no-allocate requests can be generated in two ways:

- 1) page tables that specify write-back no-allocate memory
- 2) accesses to any write-back cacheable memory when the data-cache is disabled

Note: This erratum matches bug #5124 in the ARM internal Jira database.

Conditions

- 1) A stream of read requests to write-back cacheable memory is detected by the hardware prefetch logic
- 2) The final request of this stream is to write-back read-write-allocate memory, or a load exclusive to write-back no-allocate memory (which will generate a write-back read-write
- 3) allocate request to the L2)
- 4) Hardware prefetch issues a cache line read with write-back read-write-allocate attributes
- 5) A CPU makes a request to the same line with write-back no-allocate memory attributes

Implications

If the above conditions happen enough times (minimum >8 times for a given cache set) without any of the affected lines being loaded into the data cache or clean/invalidated through software maintenance, the part could deadlock. In our hundreds of billions of simulation cycles the event occurring even a single time is very rare, but in silicon running over a long period of time without any L2 cache clean/invalidation it would be possible to generate a deadlock scenario.

Workaround

The simplest full workaround for this erratum is to disable hardware prefetch in the L2. Write 0x0000_0400 to the L2 Prefetch Control Register to accomplish this.

An alternative workaround is to prevent the LS from issuing memory requests to write-back no-allocate memory. Write-back no-allocate requests will only occur if:

- 1) the long page table format is in use and write-back-NoAllocate is specified as a memory type.
- 2) the cache is disabled, the MMU is enabled, and a data access is made to write-back memory.

If these requests can be avoided, the erratum will not occur. The operating system should not make use of WB-NoAllocate memory. When disabling the cache (during a powerdown sequence or some other cache maintenance activity) either don't access write-back cacheable memory, or disable the MMU.

766421: Strongly-Ordered/Device load or NC LDREX could return incorrect data**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

In certain situations, a strongly ordered or device load instruction, or a non-cacheable normal memory load-exclusive instruction could match multiple fill buffers and return incorrect data. This occurs because the non-speculative request no longer needs to be translated by the TLB (it has already made its request to memory). However, if its original TLB entry is lost or if it matches a different TLB entry, it will read out a random physical address (in the case of a false TLB hit) or a physical address of zero (in the case of a TLB miss). This new physical address information can cause a false hit if there happens to be a request to that physical address in flight at the same time. This will not be common.

Note: This erratum matches bugs #5139 in the ARM internal Jira database.

Conditions

- 1) MMU enabled
- 2) Non-speculative load (SO/Dev load or NC LDREX) executed, committed, causes memory read (LD1)
- 3) The TLB translation for LD1 is lost (invalidated or replaced)
- 4) Load or store to normal memory causes memory read to physical address 0x0-0x3f (OP2)
- 5) Data for both requests returns
- 6) LD1 issued down the pipe, misses TLB, reads 0x0 for physical address
- 7) LD1 Matches fill buffer entries for LD1 and OP2, merges the data

--OR--

- 1) Non-speculative load (SO/Dev load or NC LDREX) executed, committed, causes memory read (LD1)
- 2) LD1 virtual address =X, LD1 physical address=Y
- 3) The processor translation regime changes: any change to security state, ASID, VMID, MMU-enable
- 4) Load or store to normal memory causes memory read to physical address=Y (OP2)
- 5) Data for both requests returns
- 6) LD1 issued down the pipe, falsely hits TLB, reads Y for physical address
- 7) LD1 matches fill buffer entries for LD1 and OP2, merges the data

Implications

If the above combination of events occurs and satisfies the timing conditions required to trigger the erratum, the data returned to the register file for LD1 will be incorrect.

Workaround

Do all of the following:

- 1) Do not map any Normal Memory page table entry to the lowest 4k of physical memory.
- 2) Add a DMB instruction to flush all previous memory operations when making any change to the translation regime and before doing any new loads/stores/preloads in the new translation regime. Changes to the translation regime are:
 - ASID change
 - VMID change
 - MMU enable/disable
 - Security state change
 - Entering/leaving Hypervisor mode

For the security state changes, place a DMB in any secure exception handlers before any memory operations, and put a DMB after any memory operations before branching to non-secure state.

For the hypervisor entry/exit, place a DMB in any hypervisor trap or exception handlers before any memory operations, and put a DMB after any memory operations before branching back to the Guest-OS (non-secure non-hyp mode).

DRAFT

766422: Thumb store translation fault to Hypervisor may not have correct HSR Rt value**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

When a non-secure non-hypervisor memory operation instruction generates a stage2 page table translation fault, a trap to the hypervisor will be triggered. For an architecturally defined subset of instructions, the Hypervisor Syndrome Register (HSR) will have the Instruction Syndrome Valid (ISV) bit set to 1'b1, and the Rt field should reflect the source register (for stores) or destination register for loads.

On Cortex-A15, for Thumb and ThumbEE stores, the Rt value may be incorrect and should not be used, even if the ISV bit is set. All loads, and all ARM instruction set loads and stores, will have the correct Rt value if the ISV bit is set.

Note: This erratum matches bugs #5151 in the ARM internal Jira database.

Conditions

- 1) Executing in Thumb or ThumbEE mode in non-secure-non-hyp
- 2) Execute a store instruction that triggers a stage2 translation fault

Implications

The hypervisor cannot trust the Rt value for Thumb/ThumbEE stores and must load and decode the instruction manually in order to emulate the instruction. This will lead to slower hypervisor emulation support in the hypervisor when the GuestOS is executing in Thumb/ThumbEE mode.

Workaround

In the hypervisor handler, in addition to the ISV bit, the hypervisor should check the HSR.WnR bit and the mode bits in the SPSR_hyp. If the process that generated the exception was executing in Thumb/ThumbEE mode and the HSR.WnR bit is 1'b1, the hypervisor handler should consider the Rt value to be untrustworthy and decode the instruction in software.

769270: L2 Tag Latency = 3 should not be used in a system configured with Tag Slice**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

In an A15 configured with the L2 Tag Slice and programmed to an L2 Tag RAM Latency of 3 cycles, in certain situations write requests may be sent to the L2 Tag RAM as if it had a latency of 2 cycles. This could lead to corrupted Tag RAM contents and unpredictable results.

Note: This erratum matches bug #5215 in the ARM internal Jira database.

Conditions

- 1) A15 configured with the L2 Tag Slice
- 2) A15 L2 Tag RAM latency programmed to 3 cycles

Implications

On versions of A15 affected by this erratum, the above configuration should not be used to avoid data corruption and possible processor deadlock.

Workaround

The workaround for existing silicon with the L2 Tag Slice that is affected by this erratum is to program the Tag RAM latency to 2 cycles. This could affect the maximum frequency at which the part will operate.

773022: Incorrect instructions may be executed from loop buffer**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

In certain rare sequences of code, the loop buffer may deliver incorrect instructions.

NOTE: This erratum matches bug #5121 in the ARM internal Jira database.

Conditions

- 1) A15 loop buffer enabled
- 2) Rare sequence of branches executed

Implications

If the loop buffer is enabled, incorrect code may be executed.

Workaround

On affected product versions, the loop buffer should be disabled. This can be achieved by setting bit 1 of the ACTLR register to 1'b1.

774569: Watchpoint may not be taken on second half of unaligned ld/st crossing a 64-byte aligned boundary**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

An unaligned memory transaction that crosses a 64-byte aligned memory boundary is split into two requests in the A15 pipeline. Each of these requests is compared to the active watchpoints to see if a watchpoint event needs to occur.

If the address mask field in the DBGWCR is set to 'No mask' or 'Three bits masked' then the second half of such an unaligned transaction will never trigger a watchpoint.

Note: This erratum matches bug #5275 in the ARM internal Jira database.

Conditions

- 1) DBGWCR[28:24] (Mask) = 5'b00000 or 5'b00011
- 2) Unaligned memory request crossing a 64-byte boundary
- 3) DBGWVR has a watchpoint address in the upper 64-byte memory region of the request

Implications

No watchpoint will be taken.

In most cases, this will not be an issue. If a specific variable in memory is being watched, the lowest byte address of the variable should be set in the WVR and the bug will not occur. Memory requests that are moving a region of memory without regards to the underlying variable locations will tend to be aligned. The issue will only arise when the WVR location specified is not the target address of the memory transaction being issued.

Workaround

Set the DBGWCR[28:24] to 5'b00100 (4-bits masked) or higher. The watchpoint will now be taken correctly. However, Software may have to detect and ignore false triggers of the watchpoint because a watchpoint may also fire for nearby bytes in the same 64-byte aligned memory region as the desired bytes.

774769: Data corruption may occur with store streaming in a system**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

An L2 write request (A) uses a single data bank for two passes to do an ECC read/modify/write. This allows the last quad word of a cache line transaction (B) to get significantly delayed from the first quad word going through the L2 pipe. During this window the L2 buffer associated with that transaction is deallocated and reallocated to a full cache line streaming store (C) from a CPU that will allocate to the L2 cache without an external memory access. The last quad word of the original transaction then collides with the streaming store and incorrect data is written to the cache for that streaming store.

In order to get the streaming store (C) the CPU must have detected a store stream. A series of cacheable stores in a CPU triggers the store streaming mode in that CPU. The CPU then buffers up 64-byte lines of cacheable store data before sending them to the L2 cache. One of those full line writes misses the L2, but will immediately allocate to the L2 without an AXI/ACE request. This is only possible in a non-ACE system, or with streaming stores to non-shared data. Any shared data in an ACE coherent system would first need to obtain ownership of the line with an AR channel memory request.

Note: This erratum matches bug #5293 in the ARM internal Jira database.

Conditions

- 1) store streaming enabled and allocating data to the L2 cache
- 2) a series of incrementing stores is executed to an aligned 64-byte region of cacheable memory
- 3) either:
 - the memory accessed by the stores is non-shared, OR
 - the system is not coherent over ACE (BroadcastInner and BroadcastOuter pins both deasserted)

Implications

If this condition is hit, external memory may be corrupted.

Workaround

The workaround is to configure write streaming on versions of A15 affected by this erratum such that no streaming-write ever allocates into the L2 cache. This can be done by setting the no-allocate threshold to be lower than the L2-allocate threshold. Once streaming starts, all store streams will go immediately to external memory. ACTLR[28:27] (Write streaming "no-allocate" threshold) should be set to 2'b00 (12 cache lines) and ACTLR[26:25] (Write streaming "no-L1-allocate" threshold) should be set to 2'b01 (64 cache lines) or 2'b10 (128 cache lines).

777770: ESR incorrect for AdvSID HCPTR trapped inst in Hyp mode**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

The EC field of the HSR should be written to 0x7 for an HCPTR trapped coprocessor instruction while executing in HYP mode. A15 treats this as an UNDEF exception and writes 0x0 to the HSR.EC field.

Note: This erratum matches bug #5317 in the ARM internal Jira database.

Conditions

- 1) HCPTR[10] and/or HCPTR[11] set to 1'b1
- 2) A VFP/Advanced SIMD instruction is executed

Implications

A hypervisor that is trapping VFP or Advanced SIMD instructions executing in the hypervisor will see 0x0 in the EC field of the HSR. This EC value will indicate an Undefined exception. The hypervisor handler will have to load the instruction that triggered the exception and decode it to determine that the instruction was VFP or Advanced SIMD. This will take somewhat longer than determining the issue directly from the HSR register.

Workaround

If the hypervisor is running with HCPTR[10] or HCPTR[11] set and expecting to trap VFP/Advanced SIMD instructions, ensure that your Undefined exception handler loads the instruction to determine that an Advanced SIMD or VFP instruction was being executed and handle it as needed.

784420: Speculative instruction fetches with MMU disabled might not comply with architectural requirements**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

When all applicable stages of translation are disabled, an ARMv7 processor must follow some architectural rules regarding speculative fetches and the addresses to which these can be initiated. These rules avoid potential reads to read-sensitive areas. For more information about these rules see the description of "Behavior of instruction fetches when all associated MMUs are disabled" in the ARM Architecture Reference Manual, ARMv7-A and ARMv7-R edition. Cortex-A15 normally operates with both the MMU and branch prediction enabled. If the processor operates in this condition for any significant amount of time, the BTB (branch target buffer) will contain branch predictions. If both stages of translation are then disabled, but branch prediction remains enabled, these stale BTB entries can cause A15 to violate the rules for speculative fetches.

Note: This erratum matches bug #5360 in the ARM internal Jira database.

Conditions

The erratum can occur only if the following sequence of conditions is met:

- 1) MMU enabled for at least one stage of address translation
- 2) Branch prediction enabled
- 3) Branches executed
- 4) MMU disabled for all applicable stages of address translation

Note: When executing in a Non-secure PL1 or PL0 mode, for condition 1 at least one stage of address translation must be enabled, and for condition 4 both stages of address translation must be disabled. When executing in any other mode, there is only one stage of address translation, that must be enabled for condition 1 and disabled for condition 4.

Implications

If the above conditions occur, it is possible that after the MMU is disabled, speculative instruction fetches will occur to read-sensitive locations.

Workaround

Branch prediction should be disabled when the MMU is disabled after having been enabled. This should be done by clearing the appropriate Z bit in the System Control register at the same time as or just before the final stage of translation is disabled. Branch prediction should remain disabled until the MMU is enabled, or until the BTB has been flushed. On A15, the BPI* branch predictor maintenance commands will not invalidate the BTB. The BTB can be flushed by setting bit 0 of the ACTLR register, doing any instruction cache invalidate instruction (e.g. ICIALLU), and then clearing bit 0 of the ACTLR register.

784477: CTIINTACK register needs clearing each time it is set**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

The CTI contains a CTIINTACK register, which enables a trigger to be acknowledged through software, instead of using a hardware knowledge using the CTITRIGOUTACK input. The correct operation of this register is that writing a one to the bit corresponding to a trigger output will cause that trigger to be cleared, and this will not affect future triggers.

Because of this erratum, when a bit in the CTIINTACK register is set, it remains set until cleared by writing zero to the register. This causes the corresponding trigger outputs to be acknowledged immediately if they occur again, which can lead to them being missed.

The CTIINTACK register is normally used in two cases:

- To clear a debug-originated interrupt, if required by the interrupt controller.
- To clear a debug entry request generated by another processor, when cross-halting is used.

Conditions

The following conditions must occur:

- A CTI trigger output fires.
- The CTI CTIINTACK register is used to acknowledge the trigger output, by writing a one to the bit corresponding to that trigger output.
- The same trigger output fires again before the corresponding bit in the CTIINTACK register is cleared.

Implications

Trigger outputs might be missed:

- In the case of a debug-originated interrupt that uses CTIINTACK to clear the interrupt, events other than the first event might not cause an interrupt to occur.
- In the case of a cross-halting debug request, after the first time a processor halts and restarts, it might halt without halting other processors with it.

Workaround

This is a workaround for tools vendors.

When the CTIINTACK register is written with a nonzero value, it must be immediately written to again with the value zero. This prevents any future events on the corresponding trigger output from being acknowledged.

If this workaround is used, there remains a race condition, whereby a trigger output occurring between the two register writes might be lost. This is in general not significant, because the timing of trigger outputs and the timing of register writes are not highly correlated, and if the trigger output had occurred before the first register write, then it would also have been lost.

785769: Undefined exception is not generated for LDC/STC instructions which access DCC registers in User mode**Category B****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

The Debug Communication Channel (DCC) registers are accessible using MCR/MRC instructions. LDC/STC instructions provide alternate access to DCC registers DBGDTRTXint/DBGDTRRXint.

In Non-debug state when DBGDSCR.UDCCdis is set to 1, then access to DCC register using MCR/MRC and LDC/STC instructions from User mode should generate an Undefined Instruction exception. Access using MCR/MRC instructions correctly generates an Undefined Instruction exception correctly. However access using LDC/STC instructions does not generate the Undefined Instruction exception and incorrectly accesses the register.

Note: This erratum matches bug #5372 in the ARM internal Jira database.

Conditions

- 1) The processor is in Non-debug state and User mode.
- 2) DBGDSCR.UDCCdis is set to 1.
- 3) Either the Hypervisor trap to debug registers is not set (HDCR.TDA==0) or the processor is in Secure state.
- 4) LDC to DBGDTRTXint or STC to DBGDTRRXint is executed.

Implications

If LDC or STC instructions are executed in User mode, then DCC traffic between debug host and debug target from FIQ/IRQ/Supervisor/Monitor/Abort/Hypervisor/Undefined/System modes can be corrupted due to this errata.

Workaround

Tools must use MCR/MRC instructions to access Debug Communication Channel (DCC) registers from User mode, instead of LDC/STC instructions, in Non-debug state.

Alternatively, software can avoid corruption of DCC traffic by Non-secure User mode code by setting the hypervisor trap for debug register accesses (HDCR.TDA), and handling the LDC/STC instruction appropriately in hypervisor code. There is no software workaround to prevent LDC/STC instructions executing in Secure User mode from corrupting DCC traffic handled in other processor modes.

2.6. Category B (Rare)**773319: Unaligned page boundary crossing load hit by TLB invalidate can stall until next interrupt****Category B Rare****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

An unaligned cacheable writeback load instruction is being executed. It crosses a page table boundary. Both halves of the load hit in the TLB, and the data required from the lower of the pages misses in the cache. A cache line fill is issued to the L2 cache or to memory. The data required from the upper page misses the cache, but does not issue a cache line fill due to congestion.

While that fill is outstanding, a TLB invalidate is received that knocks out the TLB entry associated with the upper page, but does not affect the lower page. After that TLB invalidate is processed, the cache line fill returns. A new page table request must now be issued to the L2TLB or memory for the upper page.

A snoop request is then received by the L1 data cache targeting the just returned cache line fill. This snoop will not be processed until a the TLB request for the upper page is complete.

If the upper page miss requires a memory read and that memory read is unable to complete until the snoop completes, the processor may stall.

If another TLB maintenance operation is received by the A15 CPU just after the above snoop and before the page miss can be issued, the upper page miss may be prevented from issuing until the snoop completes and the processor may stall.

In either case, the A15 CPU will be unable to respond to the snoop or complete the load until the next branch flush, event flush, or interrupt.

Conditions

- 1) Unaligned load request that crosses a page table boundary
- 2) The lower page accessed by the load is to writeback shared memory
- 3) Both pages hit in the L1 DTLB
- 4) The required data from the lower page misses in the cache and issues a line fill
- 5) The required data from the upper page misses in the cache but does not issue a fill
- 6) A TLB invalidate command from a different CPU is received by the L1 DTLB
- 7) This TLB invalidate affects the upper page, but not the lower page (The TLB entry for the lower page must stay valid for the erratum to occur)
- 8) Data returns for the cache line fill
- 9) A snoop request is received for the cache line that just returned
- 10) Either:
 - 1) The page table walk for the just invalidated upper page requires a memory access that stalls in the memory system dependent on the above snoop, or
 - 2) A TLB maintenance operation from another CPU is received just after the above snoop command

Implications

If the above conditions occur, the CPU will be unable to complete either the unaligned load or the snoop request until the next exception (branch flush, event flush, or interrupt) occurs. No data corruption will occur. The next branch flush, fault, or interrupt will resolve the issue and execution will continue normally.

This should be very rare in real systems and should only have a temporary performance impact when it occurs. A targeted TLB invalidate that affects only one of the two pages being accessed by an active process, combined with the extremely rare timing sequence above, will make this extremely unlikely to be seen in a running system.

Workaround

There is no workaround.

2.7. Category C

763126: Three processor exclusive access livelock

Category C

Products Affected: Cortex-A15 MP Core -NEON.

Present in: r0p4

Description

In a system with three or more coherent masters that all use the ldrex/strex synchronization primitives to access a semaphore in coherent cacheable memory, there is a possibility of a livelock condition where two masters continuously attempt and fail to get the lock while the third master continuously reads the lock.

This erratum is heavily dependent on a unique set of initial conditions, and upon specific interconnect timing once the livelock has started. It is expected to be rare in a real system that the timing conditions will be hit.

An example: two cores C1 and C2 are contending for a lock using ldrex/strex, and core C3 is looping reading the same semaphore location. Once the livelock condition has started, from the perspective of C1, the sequence will look like this:

- 1) Execute ldrex, hits the cache in unique state.
- 2) External snoop takes line to shared state (triggered by C3 read).
- 3) Execute instructions to process the ldrex result and prepare the strex data.
- 4) Execute strex, hits cache shared, issues readUnique to bring in line unique.
- 5) External snoop invalidates line, clearing monitor (triggered by C2 strex that will eventually fail the monitor).
- 6) Line returns in unique state, but strex fails due to cleared monitor.
- 7) Loop back to step 1.

C1 and C2 constantly issue ReadUniques due to failing store exclusives that invalidate the line in the other core, each core causing the others strex to fail without making forward progress. No forward progress is made until/unless one of the cores stops (possibly due to an interrupt) or interconnect timing happens to allow enough time for one of them to complete.

NOTE: this erratum is describing additional limitations on exclusives loads and stores in a multi-core system including A15. There is no plan to fix this erratum on future A15 cores, as reasonable code following the ARM architecture guidelines should not be affected.

NOTE: This erratum matches bug #4637 in the ARM internal Jira database.

Conditions

- 1) One master continuously reading the location of the semaphore.
- 2) Two masters doing a ldrex/strex loop to the semaphore.
- 3) Semaphore in write-back shared memory.
- 4) Three master system (3+ core A15, or 3+ total processors in the system over ACE).

Implications

Neither C1 nor C2 will ever succeed in gaining the lock. Software could stop making progress. An interrupt to one of the cores C1/C2/C3 would likely break the livelock.

Workaround

If there are no more than two coherent masters in the system, no workaround is needed, the issue will not be seen.

The latest version of the ACE specification adds additional command types and system logic to allow processors to avoid this issue. This specification update was not available in time for A15 to take advantage of it and A15 does not implement this ACE feature. As an alternative, A15 installed hardware in each processor to detect that the load/store exclusive livelock scenario may be occurring and delay snoops for a period of time to allow the load exclusive/store exclusive loop to complete and make forward progress. With this fix, no existing code that uses ldrex/strex should need to be rewritten if it follows the ARM Architecture Reference Manual guidelines in the “A3.4 Synchronization and Semaphores” section and is not unreasonably long (>512 cycles to execute).

766174: ID_PFR0 Jazelle extension support incorrectly encoded**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

In the ID_PFR0 register, bits[11:8] identify the type of Jazelle extension support provided. Cortex-A15 currently has a value of 0x0 there (Jazelle not supported). This is a static register read by software to identify processor features that are available. The value should be 0x1 (Support for Jazelle extension, without clearing of JOSCR.CV on exception entry). Cortex-A15 does not support direct byte-code execution, but it does support the trivial implementation of the Jazelle extension.

Note: This erratum matches bugs #5126 in the ARM internal Jira database.

Conditions

Reading ID_PFR0

Implications

Software relying on that register will believe that Cortex-A15 does not support Jazelle extensions.

Workaround

Any software relying on that value to determine whether Jazelle support is available should check the MIDR to determine if it is running on a Cortex-A15. If it is running on a Cortex-A15, it can assume that the trivial Jazelle support is available.

766419: Benign op instead of UNDEFINED in certain unused opcodes**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

The ARM architecture specifies that some unallocated instruction encodings around the CLREX, DSB, DMB, and ISB instructions are illegal encodings that should take an UNDEFINED instruction exception. Cortex-A15 treats these as UNPREDICTABLE and, to avoid logic/speed-paths, treats those encodings the same as the nearby CLREX/DSB/DMB/ISB instructions.

The effect will be that, if you execute one of these unallocated encodings, you will get a CLREX or barrier instead of an UNDEFINED instruction exception.

--CLREX--

32'b1111 0101 1111 xxxx xxxx xxxx x0xx xxxx will be treated as CLREX

32'b1111 0101 x111 xxxx xxxx xxxx 10xx xxxx will be treated as CLREX

32'b1111 0101 x111 xxxx xxxx xxxx x01x xxxx will be treated as CLREX

32'b1111 0101 x111 xxxx xxxx xxxx x0x0 xxxx will be treated as CLREX

--DSB--

32'b1111 0101 1111 xxxx xxxx xxxx x100 xxxx will be treated as DSB

32'b1111 0101 x111 xxxx xxxx xxxx 1100 xxxx will be treated as DSB

--DMB--

32'b1111 0101 1111 xxxx xxxx xxxx x101 xxxx will be treated as DMB

32'b1111 0101 x111 xxxx xxxx xxxx 1101 xxxx will be treated as DMB

--ISB--

32'b1111 0101 1111 xxxx xxxx xxxx x11x xxxx will be treated as ISB

32'b1111 0101 x111 xxxx xxxx xxxx 111x xxxx will be treated as ISB

32'b1111 0101 x111 xxxx xxxx xxxx x111 xxxx will be treated as ISB

Note: This erratum matches bugs #5134 in the ARM internal Jira database.

Conditions

- 1) executing in ARM mode
- 2) execute one of the above opcodes that should trigger an UNDEFINED instruction exception

Implications

Historically in the ARM architecture these encodings were UNPREDICTABLE and have only recently been made UNDEFINED. If malicious code happens to execute one of these encodings, the effects are benign. The barriers have no architectural state impact. The CLREX will clear the local monitor, but had the encoding triggered an UNDEFINED instruction exception, the exception handler would have been required to clear the local monitor anyway.

The reason these are UNDEFINED is that the architecture may subsequently re-purpose them, thus there is an assumption that an end user should not be relying on the UNDEFINED behavior (this is what the permanently UNDEFINED space is for).

Workaround

None.

768529: Debug version of ID_PFR0 Jazelle extension support incorrectly encoded**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

In the ID_PFR0 register, bits[11:8] identify the type of Jazelle extension support provided. The correct value for Cortex-A15 should be 0x1 (Support for Jazelle extension, without clearing of JOSCR.CV on exception entry). However all r0 version (r0p0,r0p1,r0p2,r0p3) had 0x0 (Jazelle not supported). This is discussed in Erratum 766174. The encoding was corrected for r1p0.

However, the debug logic has a separate version of the ID_PFR0 register for access by external debug reads. This debug version of the ID_PFR0 register still has the 0x0 encoding for r1p0. This means that debug reads of the ID_PFR0 register will see an incorrect value.

Note: This erratum matches bugs #5158 in the ARM internal Jira database.

Conditions

Debug read of ID_PFR0

Implications

Debug software could read a value of the ID_PFR0 register that doesn't match that value seen by software executing on A15.

Workaround

Ignore the Jazelle support encoding in the debug register version of ID_PFR0.

768532: Mapping the GIC memory mapped region as cacheable can result in unpredictable behavior or a deadlock**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

If the physical address page that maps to the A15 internal GIC memory mapped registers is configured as Writeback cacheable in the page tables, a memory access to this region will generate an error and the requested line will not be placed into the L1 data cache. Unless the requesting instruction is flushed (branch or event flush) an external abort will be generated. The L1 data cache will, however, continue to evict the existing line in the data cache that was targeted for eviction by the failing cache line read. The L2 cache does not expect this eviction. This places the L2 logic in an illegal state.

If there is a snoop for the victim line while the eviction is being processed in the L2, A15 will behave unpredictably and could deadlock.

Note: This erratum matches bugs #5180 in the ARM internal Jira database.

Conditions

- 1) A15 configuration that includes the internal GIC
- 2) Page table maps the memory region assigned to the GIC memory mapped registers as WB cacheable
- 3) Memory access (load or store) made to that page
- 4) Snoop from another master (inside or outside of A15) targets the L1 victim of that memory access

Implications

Incorrect page tables mapping the GIC to cacheable could cause unpredictable behavior on A15. In general, an external data abort would be detected and the offending process can be killed. However, if the requests that caused the illegal accesses are speculative and later flushed, no abort would be reported.

If a snoop occurs to the victim line with the correct timing relative to the eviction of that line, the part could deadlock.

In addition, if the snoop collision with the victim of the illegal GIC access occurs, A15 will be operating in unvalidated space. It is possible that data in the L1 or L2 cache could become corrupted. We have not been able to find such a case, but have difficulty proving it could not happen.

Workaround

Mapping the physical address region for the GIC to cacheable memory in all page tables avoids this erratum.

A hypervisor mapping the physical memory for the GIC registers to SO or Dev in the second stage translation will prevent a Guest OS from causing this behavior.

770320: Single bit ECC error can cause cache maintenance operation to violate memory ordering**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

If a single bit ECC error occurs in the L1 data cache of A15, the L1 will make a request to the L2 logic to trigger an eviction of that line from the L1 data cache in order to correct the error. This background eviction ignores the normal memory ordering rules on cache maintenance operations. However, if there is a cache maintenance operation in flight at the time the ECC error is discovered, it is possible that this cache maintenance operation will pass a store in the machine to the same address as the cache maintenance operation. This will violate the ARM memory ordering requirements.

Note: This erratum matches bug #5223 in the ARM internal Jira database.

Conditions

- 1) Store to address A has been executed, has not yet been pushed to the L1 cache or L2 cache
- 2) DCCMVA or DCCIMVA is executed to address A after the store to address A in program order
- 3) ECC error detection logic is built into the L1 cache and enabled
- 4) Single bit ECC error is detected in the L1
- 5) The DCCMVA or DCCIMVA completes and prepares to go to the memory system after the ECC error is detected and before the ECC error implicit cache maintenance operation is issued
- 6) Software was depending on the cache being empty or clean at the end of the DCCMVA/DCCIMVA and will fail if the store data was not pushed out

Implications

The DCCMVA or DCCIMVA could pass the store to the same address. This could leave dirty data in the cache where software would expect the cache to be clean or invalid.

This erratum is likely to be rare enough to ignore. ECC errors are inherently very rare. Cache maintenance operations occurring very close to earlier stores to the same address are likely to only exist in specialized code. The odds of a store/CMO/ECC error occurring at the same point in time is very low. In general this erratum can be ignored.

When the erratum conditions do occur, the effect will be silent data corruption. A correctable error will be reported in the syndrome registers, but there will be no way to determine that the rare boundary conditions lined up to allow data corruption. However, because the alignment of events is very rare, this is not expected to add significantly to the silent error rate.

Workaround

In general, this issue should be ignored, other than to understand that affected A15 versions do not have perfect reporting or recovery from single ECC errors.

773023: Order may not be maintained between Strongly Ordered memory requests and Device memory requests on ACE/AXI**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

A15 maintains order between Strongly Ordered (SO) memory requests as required by the ARM architecture memory ordering model. This is done inside the A15 by use of internal ordering logic. On the interconnect, A15 enforces ordering by using the same ARID/AWID value for all SO memory requests from a given processor (guaranteeing read/read and write/write ordering) and by waiting for the completion of all SO and Device memory requests on one channel before issuing SO or Device requests from the same core on the other channel (guaranteeing read/write and write/read ordering).

A15 does the same for Device memory.

However, A15 uses different ARID and AWID for Device memory requests from a given CPU and Strongly Ordered memory requests from the same CPU. Due to this fact, it is possible that the an SO read from a given CPU could pass a Device read from the same CPU and arrive at a single peripheral out of order.

Conditions

- 1) A system has memory mapped peripherals larger than 4KB
- 2) Some of the pages mapped to that peripheral are mapped Strongly Ordered and some are mapped Device
- 3) Software depends upon ordering of these Strongly Ordered and Device memory requests

Implications

This is not an issue for any device that fits in one 4KB memory page, as it is only possible to have a single memory type for that page (SO/Dev aliasing is not allowed).

For larger peripherals, it is possible that Strongly Ordered or Device transactions could arrive at the peripheral out of order.

Workaround

A given peripheral device should be mapped to all Strongly Ordered or all Device memory.

774570: Fault Status bit in register DBGDSCR is implemented as sticky**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

Fault status bit[9] in debug register DBGDSCR might be improperly set for synchronous data aborts in debug state in some cases.

Note: This erratum matches bug #5278 in the ARM internal Jira database.

Conditions

- 1) Core is currently in debug state.
- 2) If either of the following occur:
 - 1) Second stage synchronous data abort occurs in Non Secure PL0 or PL1 mode when external debugger issues an instruction through DBGITR debug register, or
 - 2) Synchronous data abort other than second stage synchronous data abort occurs in Non Secure PL0 or PL1 mode when external debugger issues an instruction through DBGITR debug register with HCR.TGE bit set to 1.
- 3) FS bit in debug register DBGDSCR bit gets set due to this synchronous data abort as described above.
- 4) FS bit is not cleared by external debugger by explicitly writing 1'b0 to bit[9] of DBGDSCR.
- 5) If any of the following occur:
 - 1) Another synchronous data abort other than second stage synchronous data abort occurs in Non Secure PL0 or PL1 mode when external debugger issues another instruction through DBGITR debug register with HCR.TGE bit set to 0, or
 - 2) Another synchronous data abort occurs in Non Secure PL2 mode when external debugger issues another instruction through DBGITR debug register, or
 - 3) Another synchronous data abort occurs in Secure state when external debugger issues another instruction through DBGITR debug register.

Implications

Read of debug register DBGDSCR fault status bit[9] returns incorrect value as 1 when it should be 0.

Workaround

Whenever the external debugger reads DBGDSCR.FS bit as 1 following a synchronous data abort caused by an instruction issued through DBGITR register, the external debugger should explicitly write 1'b0 to fault status bit[9] of DBGDSCR.

774571: Sampling registers DBGCIDSR and DBGVIDSR can return incorrect value**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

Context ID sampling register DBGCIDSR and Virtualization ID sampling register DBGVIDSR can return incorrect values in some cases.

Note: This erratum matches bug #5279 in the ARM internal Jira database.

Conditions

- 1) Non invasive debug authentication status currently allows sample based profiling.
- 2) Debug software access lock status is set as indicated by SLK, bit[1] of DBGLSR.
- 3) PC sampling register DBGPCSR is read by external sampling agent through external debug interface.
- 4) Context ID or Virtualization ID has changed and ISB or exception entry or exception return has occurred and PC sampling register DBGPCSR is read by spurious software running on the current cpu or another cpu in the system using memory mapped debug interface before DBGCIDSR and DBGVIDSR registers are read by external sampling agent through external debug interface.

Implications

Reads of debug registers DBGCIDSR and/or DBGVIDSR might return a value which is not consistent with PC value sampled by external agent.

Workaround

This issue can occur only when a spurious read to DBGPCSR is made by software running on a cpu in the system. One possible workaround is to set up a translation table for the system register map such that it is not possible for software to access Cortex-A15 debug register space.

775622: Debug version of Cache Type Register (CTR) IminLine field are incorrectly encoded**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

Debug version of CTR.IminLine, register bits[3:0], are incorrectly encoded as 4'b0100 instead of 4'b0011 when IMINLN input of Cortex A15 is LOW.

Note: This erratum matches bug #5304 in the ARM internal Jira database.

Conditions

This erratum requires both of the following conditions to be met.

- 1) The IMINLN input of Cortex A15 is LOW. This input is sampled only during reset.
- 2) The debug register 833 at offset 0xD04, which is an alias of the CP15 Cache Type Register CTR, is read.

Implications

This should not create any issues with an external debugger because the actual value encoded in the register has no effect on hardware operation. CTR.IminLine is used by software to determine stride for cache maintenance operations while operating on a range of addresses. The processor uses the IMINLN input to generate this value. When IMINLN is HIGH, CTR.IminLine is encoded as 4'b0100 indicating 64 bytes. When IMINLN is LOW, CTR.IminLine is encoded as 4'b0011 indicating 32 bytes. This signal does not affect internal instruction cache line size or operation of the cache maintenance commands in hardware. Cortex-A15 hardware always uses 64 bytes as the minimum instruction cache line size.

Workaround

Ignore the CTR.IminLine value in the debug register reads of register 833, the Cache Type Register (CTR). If required, an external debugger can halt the processor and read the Cache Type Register (CTR) using a CP15 instruction to determine the value used by software.

777769: ICache parity error may not be corrected for NC code**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

If an instruction fetch to a non-cacheable page in memory gets a false hit on a line in the instruction cache due to a parity error in the instruction cache tag array, incorrect instructions may be executed.

Note: This erratum matches bug #5312 in the ARM internal Jira database.

Conditions

- 1) MMU enabled
- 2) Instruction fetch to page marked SO/Dev/Normal-Non-Cacheable
- 3) Parity error in instruction cache tag
- 4) Corrupted tag matches the physical address of the cache line being fetched

Implications

In an A15 configured with L1 parity/ECC and with parity/ECC checking enabled, in very rare circumstances a parity error can cause delivery of bad instructions while executing non-cacheable code. This is not expected to be an issue in normal systems as no normal programs will have instructions in non-cacheable memory with the MMU enabled. At boot, or any other time that the MMU is disabled, the erratum will not occur.

Workaround

Place instructions in cacheable memory whenever possible. If you must run non-cacheable code with the MMU enabled, first invalidate the instruction cache.

777771: Ordering of read accesses to the same memory location may not be ensured in the case of an unaligned load**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

An unaligned 64-byte boundary crossing load hits a cache line currently being returned from memory in one of the L1 fill buffers and returns its data to a register, executing out of order ahead of load instructions that are earlier in the program. One or both of the cache lines touched by the load are then lost from the L1 data cache (due to eviction or snoop invalidate). Another master modifies the lost cache line or lines. One of the earlier load instructions (which need not be unaligned) executes, reads the modified data, and returns it to a register. A15 should detect that the earlier load has returned more recent data (in violation of the ARM memory ordering model) and flush the younger load. In this case it does not correctly detect the hazard.

If this erratum is hit these loads are observed out of program order by the master doing the stores. This violates the ARM memory ordering model.

Note: This erratum matches bug #5328 in the ARM internal Jira database.

Conditions

- 1) Unaligned load executes out of order ahead of loads that are earlier in program order
- 2) The unaligned load touches two 64-byte aligned regions and is one of the following:
 - ldrh or ldr
 - ldm that is not 8-byte aligned
 - vld* to 32bit Si registers that is not 8-byte aligned
 - vld* to 64bit Di registers that is not 16-byte aligned
- 3) At least one of the lines read by the unaligned load is evicted from the L1 data cache (snooped or replaced)
- 4) The evicted line is modified by another master
- 5) One of the earlier loads now executes, brings in the modified line, and returns one of the modified bytes

Implications

In general, when one master is modifying a memory location and another is reading from it, the accesses will be synchronized with the appropriate use of semaphores or locks such that the write to memory is complete before the read is executed. Such code will not be affected by this erratum.

However, certain lock-free synchronization techniques depend on the memory order property that earlier loads will not return more recent data for a given byte in memory than later loads. If an unaligned 64-byte boundary crossing variable is used in lock-free programming it may not work correctly.

This is expected only to happen in hand-written assembly code. In a compiler, variables that are used for lock-free synchronization must be marked as 'volatile' to avoid normal compiler reordering. Current compilers will not place volatile variables in an unaligned memory location and will not generate code that will be affected by this erratum.

Workaround

Whenever possible, use aligned memory locations for any memory reads that will rely on in order observation by other masters. Volatile variables in the ARM and GNU compilers will meet this requirement for volatile variables.

If there is an unaligned memory read (as defined in condition 2 above) that must be observed in order with earlier reads, place a DMB instruction before it. This will correctly enforce the observation order.

777772: Device LDM may stall if memory type changed asynchronously from Device to Normal**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

A LDM/VLDM or extension register element load instruction begins execution. This load will be broken up into multiple micro-ops of 8 or 16 bytes. The first few load micro-ops hit the TLB and return translations with the memory type Strongly Ordered or Device. Before the rest of the micro-ops in the load instruction can complete, a TLB invalidate instruction from another master removes that TLB entry. The LDM therefore triggers a new tablewalk. This tablewalk returns a translation with a memory type of Normal. The CPU will be unable to complete the load. The CPU will stall until the next interrupt. If this occurs in the highest level of privilege with interrupts disabled, the CPU may deadlock.

Note: This erratum matches bug #5329 in the ARM internal Jira database.

Conditions

- 1) A load instruction of greater than 8-bytes is executed
- 2) The current translation for the targeted page in memory is Strongly Ordered or Device
- 3) A TLB invalidate from another CPU is received in the middle of executing the large load instruction
- 4) The tablewalk for the invalidated page returns a Normal memory page (NC, write-through, or write-back)

Implications

If all the conditions occur, the CPU will stall until the next interrupt.

This is not expected to be an issue in any real systems. An OS will not remap a page from an SO/Device type to normal memory while an active process is accessing that memory.

If in a rare case the OS does shift an actively accessed page from Strongly Ordered/Device to a Normal (NC/WT/WB) memory type, it is possible that the thread accessing the page may stall until the next interrupt. As this is expected to occur extremely infrequently or never, the impact will be low.

Workaround

When remapping a virtual address page from Strongly Ordered or Device to a Normal memory type, the OS should first remap the page from the original memory type to an invalid page, execute a DSB to ensure it is complete, and then remap the page as Normal memory. This will avoid this issue.

777774: DCCMVAC/DCCSW colliding with ReadOnce from ACE could cause data corruption**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

A cache maintenance instruction (DCCSW to the level 2 cache or DCCMVAC) is executed to a line that is dirty in the A15 L1 data cache or L2 unified cache. In a coherent ACE system, this will cause A15 to generate an ACE WriteClean operation pushing out the dirty data while retaining the line in the cache in UniqueClean state. The cache state is updated and the dirty data is placed in a buffer to be sent over the ACE write channel.

A ReadOnce command for the same line is then received over the AC snoop interface. The L2 cache detects that the WriteClean is in the buffer and sends the dirty data out over the CDATA channel in response to the snoop, marking the data as dirty. A15 has not completed the snoop and the line is in UniqueClean state. The interconnect does not yet finish writing the dirty data back to memory, however.

A write from an A15 CPU or ACP then modifies that line in the L2 cache, changing its state to UniqueDirty. The line is then replaced in the L2 cache and a WriteBack is issued on the ACE write channel.

It is possible that this WriteBack will arrive at memory before the memory update generated from the CDATA snoop response. If this occurs, the most recent WriteBack data will be overwritten by the older CDATA data. This will lead to stale data in memory.

To avoid the above case, the ACE protocol states that a UniqueDirty line can't transition to UniqueClean in response to a ReadOnce. A15 violates this restriction.

Note: This erratum matches bug #5334 in the ARM internal Jira database.

Conditions

- 1) Cache line in Unique Dirty state (Modified) in A15 L1 data or L2 cache
- 2) DCCSW or DCCMVAC instruction executed to that line
- 3) Before the WriteClean command is issued to the bus, a ReadOnce occurs for that line
- 4) The dirty data is sent on the CDATA channel in response to the snoop, but stalls in the system
- 5) That cache line is further modified by A15
- 6) That cache line is cleaned or replaced in the A15 L2 and generates a WriteBack or WriteClean
- 7) The WriteBack or WriteClean completes to memory before the stalled CDATA channel response above completes

Implications

If the above conditions are met, the latest data from A15 to that line will be overwritten by the stale data in the CDATA response. This is expected to be very uncommon or impossible in real systems, however.

A ReadOnce command will be received by A15 if a non-caching master is making a read request for that line. This would tend to imply that cache coherence for that memory location is being handled by the ACE hardware coherence mechanism.

If a DCCMVAC instruction (Data Cache Clean by MVA to point of coherence) is executed for a given memory location, there is software cache maintenance occurring. Typically software would be pushing data out of the L1/L2 caches so it can be seen by an external master that is not participating in ACE hardware coherence.

A DCCSW instruction is generally used only as part of loop cleaning all dirty data out of the caches in preparation for powering down A15.

This erratum can only occur if a ReadOnce occurs concurrent with one of the above software clean instructions. In the case of the DCCSW power down sequence, it is possible that a non-caching master would generate a ReadOnce while a DCCSW is executing. However, even if this occurs, during the core powerdown operation software will typically have the caches disabled and will be making no further modifications to those lines. There will be no second modification of the line to generate the second WriteBack/WriteClean required by the race condition.

In the case of software communicating with an external peripheral through a cached memory buffer, this erratum is expected to be rare for two reasons. First, generally only software or hardware coherence will be used for a given memory location, not both. This means that getting a ReadOnce and a DCCMVAC to the same location concurrently

would not be expected. Second, software will generally modify a section of memory and then clean that memory from the caches to be used by the non-hardware-coherent peripheral. While that peripheral is using the data, it will be unusual for software on the A15 to further modify that memory location. Because of this, the second WriteBack/WriteClean required to generate the race condition will be unlikely.

DCCMVAU (clean to point of unification), often used for self-modifying code, will not trigger this erratum.

Workaround

Do not execute DCCSW or DCCMVAC instructions to cache lines that might be accessed by an ACE ReadOnce command while A15 is actively modifying it. Replace any DCCSW that are not part of a powerdown sequence with DCCISW (clean/invalidate). If cache clean by MVA (DCCMVAC) must be done to a region of memory being accessed by a coherent non-caching peripheral, do a clean/invalidate by MVA (DCCIMVAC) instead.

DRAFT

780121: PTM might not acknowledge a trace flush request when cpu is in WFI.**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

When a trace sink requests a flush of a trace source by asserting AFVALID, the trace source normally acknowledges the flush request by asserting AFREADY after all the trace in FIFO has been output.

Under certain conditions, the PTM might not acknowledge a trace flush request even after all the trace has been output.

Note: This erratum matches bug #5346 in the ARM internal Jira database.

Conditions

The following sequence must occur:

- 1) A processor power on reset or debug trace reset request
- 2) The PTM is enabled
- 3) The PTM is later disabled
- 4) A WFI or WFE instruction is executed
- 5) The processor stops its internal clocks
- 6) The processor is still powered up

In the above scenario after step 4, PTM will de assert AFREADYM temporarily for few processor clock cycles and few trace clock cycles. If the processor clocks are stopped before AFREADYM is asserted again, then any new flushes will not be acknowledged until the processor clocks are restarted.

AFREADYM will be correctly asserted during WFI in any of the following cases:

- The processor is powered down and clamps are activated
- The PTM is reset
- The PTM was never enabled out of trace reset
- The PTM is not disabled
- The PTM is disabled during WFI

Implications

If a trace sink initiates a flush request and then stops on the flush completion, then due to this erratum the flush never completes and the trace sink will not stop.

If the PTM is connected to a CoreSight trace funnel, then if a flush request is initiated the funnel will wait for the all trace sources to acknowledge the flush before continuing. This erratum might cause the funnel to stop accepting new trace from all other trace sources, preventing any further trace capture.

Workaround

There are two workarounds:

- If the PTM is connected to a trace funnel, after disabling the PTM, the trace analyzer must disable the corresponding trace slave port in trace funnel. This will make the trace funnel acknowledge the flush immediately.
- Instead of disabling the PTM, the PTM should be configured to generate no more trace, but must remain enabled.

784419: An unaligned load crossing a 4K boundary between SO/Dev memory and WB memory can stall the core until the next interrupt**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

If an unaligned load crosses a 4k page boundary between two pages where the lower page is mapped to Strongly-ordered or Device memory and the upper page is mapped to Normal, Write-Back Cacheable memory, in certain rare cases the core may stall until the next interrupt. A memory transaction that crosses a boundary between these page types is architecturally UNPREDICTABLE and should not occur in any real code. However, if this were to occur at the highest privilege level with interrupts disabled, it could deadlock that CPU.

Note: This erratum matches bug #5355 in the ARM internal Jira database.

Conditions

- 1) A load instruction is executed that accesses bytes from two different 4k pages
- 2) The lower page is Device or Strongly-ordered memory type
- 3) The upper page is Normal Write-Back Cacheable memory, that is also Read-Allocate, Write-Allocate, or both
- 4) The load instruction is one of the following:
 - LDRH or LDR
 - an LDM that is not 8-byte aligned
 - a VLD* to 32bit Sd registers that is not 8-byte aligned
 - a VLD* to 64bit Dd registers that is not 16-byte aligned

Implications

If the above conditions occur, it is possible that in rare cases the core will stall until the next interrupt. If this occurs in a situation where no interrupt will occur, the CPU deadlocks. Examples of where a deadlock might occur are if there is no timer interrupt in the system, or the conditions occur at the highest privilege level with interrupts disabled.

Workaround

Do not execute loads that cross between Strongly-ordered or Device memory and Normal memory pages. Architecturally, such a load is UNPREDICTABLE.

784469: CTI Authentication Status register is incorrect**Category C****Products Affected: Cortex-A15 MP Core -NEON.****Present in: r0p4****Description**

The AUTHSTATUS register is a read-only register in the CTI that reports the debug level supported by the CTI and the current status of the debug level.

The CoreSight Architecture Specification specifies bits [3:0] in the AUTHSTATUS register as below:

- [3:2] Non-Secure Non-Invasive Debug
- [1:0] Non-Secure Invasive Debug

For each of these fields, the value of the status bits as returned by the CTI and their meanings are specified as below:

Value Description

2'b10 Functionality disabled

2'b11 Functionality enabled

In the CTI each pair of bits ([3:2] and [1:0]) in the AUTHSTATUS register currently read:

- When functionality is disabled - 2'b01
- but should read (as per table above):
- When functionality is disabled - 2'b10

The bits are swapped.

Condition 1

AUTHSTATUS[1:0] - Non-secure Invasive Debug

- DBGEN input to the CTI is LOW
- AUTHSTATUS register is read

Condition 2

AUTHSTATUS[3:2] - Non-secure non-Invasive Debug

- NIDEN input to the CTI is LOW
- DBGEN input to the CTI is LOW
- AUTHSTATUS register is read

Implications

The status of the debug level supported by the CTI as returned by the AUTHSTATUS register read is incorrect. The masking of trigger inputs and outputs using DBGEN and NIDEN is not affected by this erratum. The return of an incorrect value might lead to incorrect operation of debug tools.

Workaround

This is a workaround for users and tools vendors. When reading the AUTHSTATUS register, swap the bits in the affected fields and interpret the read data accordingly.